Scenario Planning for Sea Level Rise via Reinforcement Learning

Salman S. Shuvo, Yasin Yilmaz, Alan Bush, Mark Hafen

***Abstract*—Climate change and sea level rise impacts will affect coastal communities with multiple threats, including increased frequency of compound events, such as storm surge combined with heavy precipitation. Accurately modeling how the stakeholders, such as governments and residents, may respond to sea level rise scenarios (i.e., scenario planning) can assist in the creation of policies tailored to local impacts and resilience strategies. In this paper, our contributions are twofold. Firstly, considering a single-agent model for government, we numerically show that the government's policy on infrastructure improvement should be based on the observed sea levels rather than the observed cost from nature. The latter refers to the straightforward policy that any responsive (but not proactive) government would follow. Through a reinforcement learning algorithm based on a Markov decision process model we show that the precautionary measures, (i.e., infrastructure improvements triggered by the sea levels) are more effective in decreasing the expected cost than the aftermath measures triggered by the cost from nature. Secondly, to generate different scenarios we consider several sea level rise projections by NOAA, and model different government and resident prototypes using cooperation indices in terms of being responsive to the sea level rise problem. We present a reinforcement learning algorithm to generate simulations for a set of scenarios defined by the NOAA projections and cooperation indices.**

I. Introduction

Global climate change and its impacts, in terms of sea level rise, have been extensively documented, analyzed and projected [1]–[3]. Climate change and sea level rise impacts will affect coastal communities with multiple threats, including increased frequency of compound events - such as storm surge combined with heavy precipitation [4]. This exacerbates social vulnerability, particularly in underserved communities [5], [6]; stresses coastal ecosystems [7], [8]; and impacts local economies by affecting property values, the tax base, and the cost of insurance, among other factors [9]. Because of sea level rise, coastal communities are vulnerable to many of these impacts, and must build the adaptive capacity and resilience frameworks to respond to these stressors through effective decision support and planning [10].

Overall, better information is needed for governments, planners, coastal managers, and personnel in a variety of agencies for effective communication, decision making and adaptation planning [10], [11]. This requires the participation of key actors to communicate the science, the variability, and the risk of various scenarios to stakeholders [12]. Accurately modeling how these agents may respond to sea level rise scenarios can assist in the creation of policies tailored to local impacts and resilience strategies, and requires a variety of community engagement and planning tools, including scenario planning [13]–[15].

This work is supported by NSF award no. 1737598. Authors are with the University of South Florida, Tampa, FL 33620, USA. (e-mail: yasiny@usf.edu).

Reinforcement learning (RL) provides a suitable theoretical framework for generating agent-based scenarios [16]. The RL agent interacts with the environment by taking an action at each time and receiving a cost/reward from the environment in return. The objective of the agent is to minimize/maximize an expected sum of costs/rewards over time by choosing optimal actions from an action set. At each time, as a result of agent's action, the system moves to a new state according to a probability distribution. The optimal policy for deciding on actions maps system states to actions, i.e., determines which action to take in which state [17].

In this paper, we consider a city setup with government as the decision maker (i.e., RL agent), and nature and residents as the environment which the agent interacts with. Our contributions are twofold. Firstly, considering a single-agent model for government, we numerically show that a rational government's policy on infrastructure improvement should be based on the observed sea levels rather than the observed cost from nature. The latter is the straightforward policy that any responsive (but not proactive) government would follow. Through simulations we show that the precautionary measures, (i.e., infrastructure improvements triggered by the sea levels) are more effective in decreasing the total cost than the aftermath measures, (i.e., infrastructure improvements triggered by the cost from nature). Secondly, to generate different scenarios we consider several sea level projections by National Oceanic and Atmospheric Administration (NOAA), and model different government and resident prototypes using cooperation indices in terms of being responsive to the sea level rise problem. The optimum policy depends on these cooperation indices, and can be found using reinforcement learning techniques. We present a reinforcement learning algorithm to generate simulations for a set of scenarios defined by the NOAA projections and cooperation indices.

The remainder of the paper is organized as follows. In Section II, the proposed RL model is explained. A reinforcement learning algorithm for finding the optimal policy is presented in Section III. Scenario simulations are given in Section IV, and the paper is concluded in V.

II. RL Problem Formulation

We investigate the problem of when to invest in infrastructure improvement against sea level rise, e.g., storm water drainage system, sea wall, levee, etc. Hence, at every time step, e.g., a year, government makes a decision $x_n = 1$ (invest) or $x_n = 0$ (no invest) for infrastructure improvement.

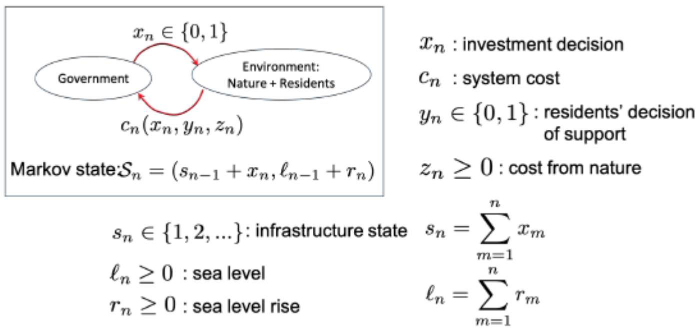

Fig. 1. Proposed MDP model.

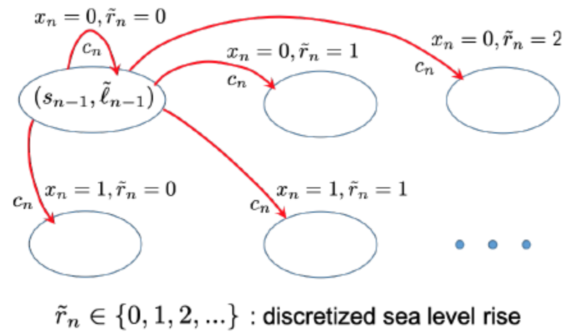

Fig. 2. MDP state transition diagram.

Accordingly, the current state of the city infrastructure against sea level rise is given by $s_n = \sum_{m=1}^{n} x_m = s_{n-1} + x_n$. Denoting the sea level rise in time interval n with $r_n \geq 0$ the sea level is similarly given by $\ell_n = \sum_{m=1}^{n} r_m = \ell_{n-1} + r_n$. We define the system state $\mathcal{S} = (s_n, \ell_n)$ as the pair of infrastructure state and sea level, which clearly satisfies the Markov property: $\mathsf{P}(\mathcal{S}_n | \mathcal{S}_{n-1}, \ldots, \mathcal{S}_0) = \mathsf{P}(\mathcal{S}_n | \mathcal{S}_{n-1})$.

Let us denote the cost from nature at each time step with z_n, e.g., the cost of flooding, storm surge, hurricane, etc. Also denote with y_n the response of residents at each time step. In this work, we use a binary response $y_n \in \{0, 1\}$ for the residents considering the residents' decision to support ($y_n = 1$) or not ($y_n = 0$) the government's investment for handling the sea level rise problem, e.g., by paying an extra tax. In this work, we investigate government's decision under sea level rise scenarios using a single-agent RL model. In this context, agent refers to the government, and environment refers to the nature and residents together. The responses x_n, y_n, z_n of government, residents, and nature, respectively, together define the agent's cost $c_n = (2 - y_n)x_n + z_n$ at each time n. This cost is normalized by taking the investment cost $(2 - y_n)x_n$ a unit cost with the resident support (two units without the resident support). The second component z_n is the normalized cost from nature with respect to the investment cost. Finally, the cumulative cost function for the agent is given by

$$C_N = \sum_{n=0}^{N} a_g^n [(2 - y_n)x_n + z_n], \tag{1}$$

which represents the total discounted cost in N time steps from now. The discount factor $a_g \in (0, 1)$ determines the weight (i.e., importance) of future costs in agent's decisions. Aside from being a standard parameter in RL cost function, a_g has an important contextual meaning in this work. It represents how much the government values the future costs due to the sea level rise problem in its decision making process. Hence, we call a_g *government's cooperation index*.

The agent's objective is to minimize $\mathsf{E}[C_N]$ by choosing its actions $\{x_n\}$ over time. This defines a Markov Decision Process (MDP), as summarized in Fig. 1. Every time agent takes an action x_n, environment reacts to that by incurring a cost c_n. Considering discretized sea level rise values $\tilde{r}_n \in \{0, 1, 2, \ldots\}$ the state transition diagram is given by Fig. 2. In our problem, environment consists of nature and residents.

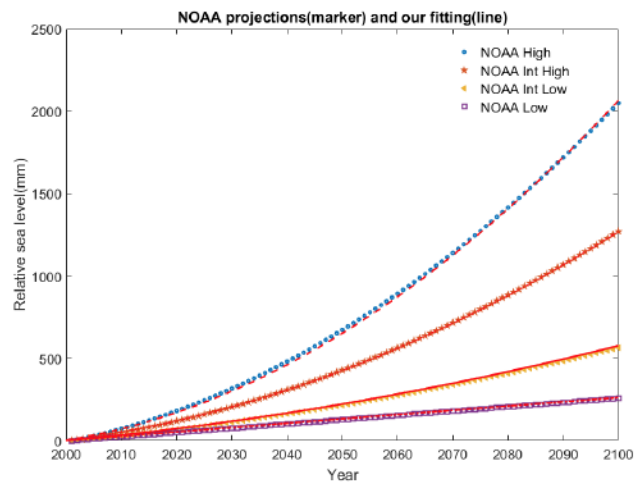

Fig. 3. NOAA projections (solid lines) for St. Petersburg, FL [18], and our curve fitting (dashed lines) for the expectation of Gamma distribution.

We next discuss suitable models for them to complete the proposed MDP model.

To model nature's cost we start with four different NOAA projections for sea level rise in St. Petersburg, FL [18], given by Fig. 3. Since these are some expected levels, in our simulations we account for uncertainty by modeling the sea level rise variable using a Gamma distribution, $r_n \sim \mathrm{Gamma}(\alpha, \beta)$. We set $\beta = 0.5$ and α to match $\mathsf{E}[\ell_n]$ with the NOAA projection curves. The successful curve fitting shown Fig. 3 is obtained by setting $\alpha = 5.2$ fixed for the low projection; by increasing α from 6.3 to 16.596 with 0.104 increments for the intermediate-low projection; by increasing α from 9.5 to 42.17 with 0.33 increments for the intermediate-high projection; and by increasing α from 14 to 69.44 with 0.56 increments for the high projection. Then, we model nature's cost z_n using the Generalized Pareto distribution, which is commonly used to model catastrophic losses [19]. The parameter settings for z_n are given as

$$\begin{aligned} z_n &\sim \mathrm{GeneralizedPareto}(k, \sigma_n, \theta) \\ k &= -0.001,\ \theta = 0,\ \sigma_n = \frac{\eta(\ell_{n-1})^a}{(s_{n-1})^b}. \end{aligned} \tag{2}$$

Through the parameters η, a, b we control the impact of most recent sea level ℓ_{n-1} over nature's cost z_n relative to the most recent infrastructure state $s_{n\;1}$.

Residents' decision y_n is modeled using Bernoulli distribution with probability parameter designed through logistic sigmoid function:

$$y_n \sim \text{Bernoulli}(p_n), \; p_n = \sigma(q_n) = \frac{1}{1 + e^{\;(q_n \;\; q_0)}}$$
$$q_n = \sum_{m=1}^{n\;1} a_r^{n\;\;m} x_m z_m, \tag{3}$$

where the score q_n reflects the willingness of residents to share government's investment costs based on the cooperation index $a_r \in (0, 1)$. It takes a high value, and yields a high probability of support if the residents' cooperation index is high ($a_r \approx 1$), nature's cost has been serious and the government has been responsive especially recently. If at least one of these conditions do not exist, then q_n tends to get smaller values, decreasing the probability of support. An average value for q_n is used for the sigmoid's midpoint q_0.

III. Finding the Optimal Policy

RL provides a data-driven solution to MDP problems. It typically updates a value function $V(s_n, \ell_n) = \min_{x_n} \mathsf{E}[C_N | x_n]$ in an iterative way based on the Bellman equation

$$\begin{aligned} V(s_{n\;1}, \ell_{n\;1}) &= \min_{x_n} \mathsf{E}[c_n + V(s_n, \ell_n) | x_n, y_n] \\ &= \min\{z_n + a_g V(s_{n\;1}, \ell_n), \\ &\quad 2 - y_n + z_n + a_g V(s_{n\;1} + 1, \ell_n)\}. \end{aligned} \tag{4}$$

The value function $V(s_n, \ell_n)$ defines the optimal policy:

$$\begin{cases} x_n = 1 & \text{if } 2 - y_n + a_g V(s_{n\;1} + 1, \ell_n) < a_g V(s_{n\;1}, \ell_n), \\ x_n = 0 & \text{otherwise.} \end{cases} \tag{5}$$

Although there are several RL algorithms, in general the RL approach learns the value function by experiencing actions and the corresponding costs. In Algorithm 1, we provide an RL algorithm based on Monte-Carlo simulations to learn the optimal government policy on infrastructure investment actions.

Algorithm 1 RL algorithm for learning optimum policy

- *Input:* a_g, a_r, *Returns*(s, ℓ): an array to save states' returns in all iterations;
- *Initialize:* $V(s, \ell) \leftarrow 0, \forall s, \ell$;
- **for** iteration $= 0, 1, 2, \ldots$ **do**
  - Generate an episode: Take actions using (5) for N steps
  - $G(s, \ell) \leftarrow$ sum of discounted rewards from (s, ℓ) till termination for all states appearing in the episode;
  - Append $G(s, \ell)$ to *Returns*(s, ℓ);
  - $V(s, \ell) \leftarrow$ average(*Returns*(s, ℓ));
  - **if** $V(s, \ell)$ converges for all s, ℓ **then**
    - **break**
  - **end if**
- **end for**

Algorithm 1 runs several episodes to iteratively compute the value function $V(s, \ell)$ for all feasible states. Each episode is a Monte-Carlo simulation in which several states are visited according to the current policy defined by the current value function. At the end of each episode the values of visited states are updated using the return, i.e., total discounted cost from a state until termination, from these states. After the state values converge, the final state values are used for generating scenario simulations, as described in the next section. The convergence of Algorithm 1 is illustrated in Fig. 4.

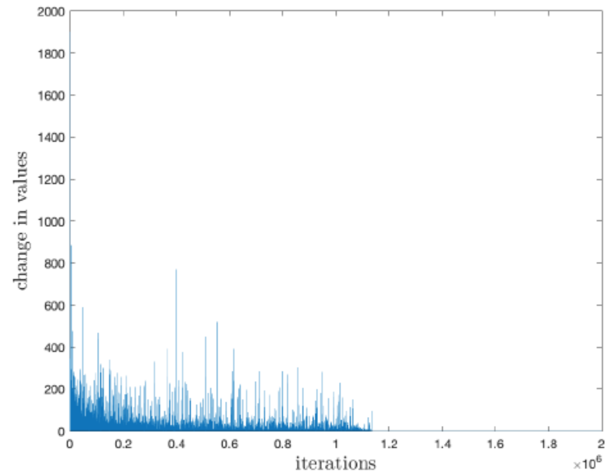

Fig. 4. Convergence of value function $V(s, \ell)$ in Algorithm 1 for the NOAA's intermediate-low projection.

IV. Scenario Simulations

In this section, we present simulation results for several sea level rise scenarios. To match the cost due to hurricanes each year that is reported in [20], we choose the parameters $\eta = 2, a = 0.4, b = 0.5$ in (2) for the nature model. We obtain different scenarios by varying the cooperation indices a_g and a_r for the government and residents, respectively, and by considering different NOAA projections for sea level rise.

Fig. 4 shows the convergence of the state values $V(s, \ell)$ in Algorithm 1 considering the NOAA int-low projection. For each scenario, once the converged state values are found, the resultant optimal policy is used to assess the cost of the RL-based government . Note that the RL-based government is proactive in dealing with the sea level rise problem as it monitors the sea level state together with the infrastructure state, and takes precautionary measures by improving the infrastructure whenever the expected future cost of not improving exceeds the improvement cost. Consider a reactive/responsive real-world government that follows a straightforward policy by improving infrastructure after experiencing a significant cost from the nature.

In Fig. 5, we compare the optimal RL policy with this straightforward policy in terms of average total cost in 100 years for the low and high NOAA sea level projections. It is seen that with the same number of investments on average the proactive policy that acts according to the sea level and infrastructure state instead of the ultimate cost from nature greatly reduces the total cost for the government.

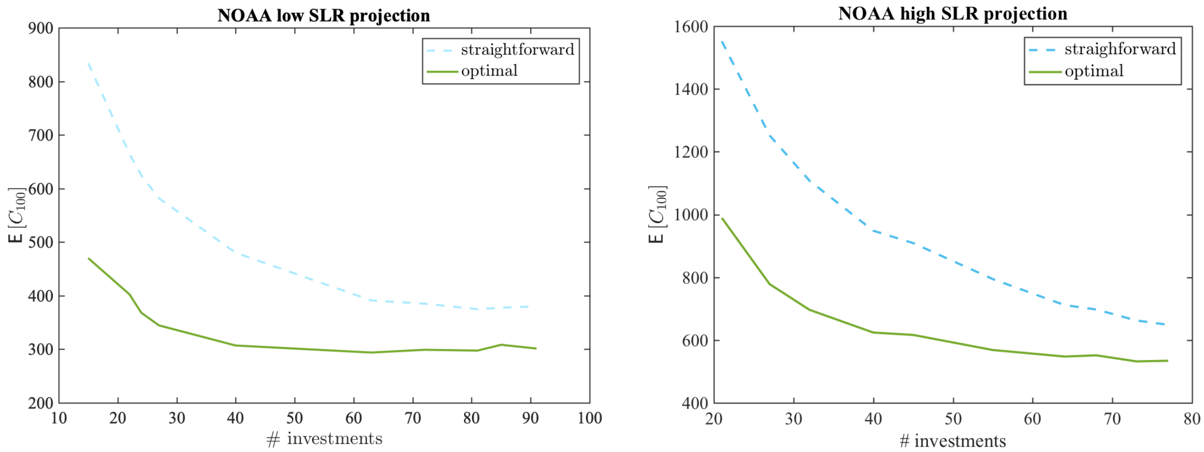

Fig. 5. Optimum policy vs. straightforward policy in terms of average total cost in 100 years for the low (left) and high (right) NOAA sea level rise projections.

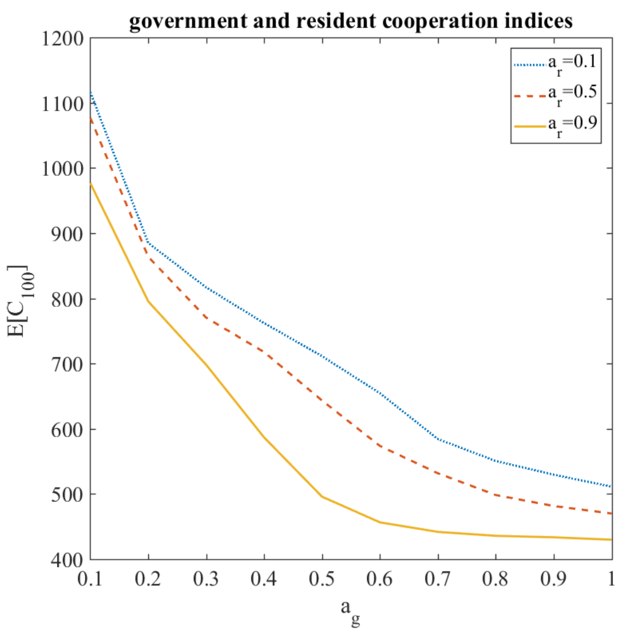

Fig. 6. Average total cost as a function of government cooperation index a_g for different resident cooperation indices a_r for NOAA high SLR projection.

Finally, in Fig. 6, we analyze the effect of cooperation indices. As expected, the average total cost decreases with growing cooperation index for both the government (a_g) and residents (a_r). The cost is more than doubled if both the government and residents are not cooperative ($a_g = a_r = 0.1$) compared to the full cooperative case ($a_g = 1, a_r = 0.9$).

V. Conclusions

In this paper, we presented a proactive government model for the sea level rise problem in a city environment considering the impacts from nature and residents. The proactive government, which learns the optimal infrastructure investment policy (yes or no at each time step) through reinforcement learning to minimize the expected economic cost over time, monitors the sea level state together with the infrastructure state, and makes an infrastructure investment to alleviate the effects of sea level rise problem whenever the expected future cost of no investment exceeds the immediate cost of investment. This proactive strategy was shown to greatly outperform the straightforward investment policy which improves the infrastructure in the aftermath of a serious economic cost from the nature. We also demonstrated that the average total cost can be significantly reduced as the government and residents become more cooperative in addressing the sea level rise problem. For the sea level rise amounts over time, different NOAA projections are considered.

References

[1] A. C. Hine, D. P. Chambers, T. D. Clayton, M. R. Hafen, and G. T. Mitchum, *Sea Level Rise in Florida: Science, Impacts, and Options*. University Press of Florida, 2016. [Online]. Available: https://books.google.com/books?id=dSTNjwEACAAJ

[2] A. H. Sallenger Jr, K. S. Doran, and P. A. Howd, “Hotspot of accelerated sea-level rise on the Atlantic coast of North America,” *Nature Climate Change*, vol. 2, no. 12, pp. 884–888, 2012.

[3] C. J. Hapke, O. Brenner, R. Hehre, and B. J. Reynolds, “Assessing hazards along our Nation’s coasts,” US Geological Survey, Tech. Rep., 2013.

[4] T. Wahl, S. Jain, J. Bender, S. D. Meyers, and M. E. Luther, “Increasing risk of compound flooding from storm surge and rainfall for major US cities,” *Nature Climate Change*, vol. 5, no. 12, p. 1093, 2015.

[5] S. B. Kashem, B. Wilson, and S. Van Zandt, “Planning for climate adaptation: Evaluating the changing patterns of social vulnerability and adaptation challenges in three coastal cities,” *Journal of Planning Education and Research*, vol. 36, no. 3, pp. 304–318, 2016.

[6] G. Schrock, E. M. Bassett, and J. Green, “Pursuing equity and justice in a changing climate: Assessing equity in local climate and sustainability plans in US cities,” *Journal of Planning Education and Research*, vol. 35, no. 3, pp. 282–295, 2015.

[7] R. J. Nicholls, “Planning for the impacts of sea level rise,” *Oceanography*, vol. 24, no. 2, pp. 144–157, 2011.

[8] T. Wahl, F. M. Calafat, and M. E. Luther, “Rapid changes in the seasonal sea level cycle along the US Gulf coast from the late 20th century,” *Geophysical Research Letters*, vol. 41, no. 2, pp. 491–498, 2014.

[9] X. Fu, J. Song, B. Sun, and Z.-R. Peng, “Living on the edge: Estimating the economic cost of sea level rise on coastal real estate in the Tampa Bay region, Florida,” *Ocean & Coastal Management*, vol. 133, pp. 11–17, 2016.

[10] T. Beatley, *Planning for coastal resilience: best practices for calamitous times*. Island Press, 2012.

[11] X. Fu, M. Gomaa, Y. Deng, and Z.-R. Peng, “Adaptation planning for sea level rise: A study of US coastal cities,” *Journal of Environmental Planning and Management*, vol. 60, no. 2, pp. 249–265, 2017.

[12] J. Tribbia and S. C. Moser, "More than information: what coastal managers need to plan for climate change," *Environmental science & policy*, vol. 11, no. 4, pp. 315–328, 2008.
[13] P. R. Berke and M. R. Stevens, "Land use planning for climate adaptation: Theory and practice," *Journal of Planning Education and Research*, vol. 36, no. 3, pp. 283–289, 2016.
[14] A. Drogoul, "Agent-based modeling for multidisciplinary and participatory approaches to climate change adaptation planning," in *Proceedings of RFCC-2015 Workshop, AIT*, 2015.
[15] R. Potts, L. Jacka, and L. H. Yee, "Can we Catch em All'? An exploration of the nexus between augmented reality games, urban planning and urban design," *Journal of Urban Design*, vol. 22, no. 6, pp. 866–880, 2017.
[16] L. Panait and S. Luke, "Cooperative multi-agent learning: The state of the art," *Autonomous agents and multi-agent systems*, vol. 11, no. 3, pp. 387–434, 2005.
[17] R. S. Sutton and A. G. Barto, *Reinforcement learning: An introduction*. MIT press, 2018.
[18] T. B. C. S. A. Panel, *Recommended Projection of Sea Level Rise in the Tampa Bay Region*. [Online]. Available: http://www.tbrpc.org/wp-content/uploads/2019/01/SLR_Recommendation_OBRC_10-9-15.pdf
[19] C. Kleiber and S. Kotz, *Statistical size distributions in economics and actuarial sciences*. John Wiley & Sons, 2003, vol. 470.
[20] J. Weinkle, C. Landsea, D. Collins, R. Musulin, R. P. Crompton, P. J. Klotzbach, and R. Pielke, "Normalized hurricane damage in the continental United States 1900–2017," *Nature Sustainability*, vol. 1, no. 12, p. 808, 2018.